

AD-A141 541

ON THE EQUIVALENCE OF LOGICAL DATABASES(U) STANFORD
UNIV CA DEPT OF COMPUTER SCIENCE G M KUPER ET AL. 1984
AFOSR-TR-84-0360 AFOSR-80-0212

1/1


UNCLASSIFIED

F/G 12/1

NL



END
DATE
FILMED
7-84
DTIC



MICROCOPY RESOLUTION TEST CHART
NATIONAL BUREAU OF STANDARDS-1963-A

UNCLASSIFIED
SECURITY CLASSIFICATION OF THIS PAGE

REPORT DOCUMENTATION PAGE

1a. REPORT SECURITY CLASSIFICATION UNCLASSIFIED		1b. RESTRICTIVE MARKINGS	
2a. SECURITY CLASSIFICATION AUTHORITY		3. DISTRIBUTION/AVAILABILITY OF REPORT Approved for public release; distribution unlimited.	
2b. DECLASSIFICATION/DOWNGRADING SCHEDULE			
4. PERFORMING ORGANIZATION REPORT NUMBER(S)		5. MONITORING ORGANIZATION REPORT NUMBER(S) AFOSR-TR- 34-0360	
6a. NAME OF PERFORMING ORGANIZATION Stanford University	6b. OFFICE SYMBOL (If applicable)	7a. NAME OF MONITORING ORGANIZATION Air Force Office of Scientific Research	
6c. ADDRESS (City, State and ZIP Code) Computer Science Department Stanford CA 94305		7b. ADDRESS (City, State and ZIP Code) Directorate of Mathematical & Information Sciences, Bolling AFB DC 20332	
8a. NAME OF FUNDING/SPONSORING ORGANIZATION AFOSR	8b. OFFICE SYMBOL (If applicable) NM	9. PROCUREMENT INSTRUMENT IDENTIFICATION NUMBER AFOSR-80-0212	
8c. ADDRESS (City, State and ZIP Code) Bolling AFB DC 20332		PROGRAM ELEMENT NO. 61102F	PROJECT NO. 2304
		TASK NO. A2	WORK UNIT NO.
11. TITLE (Include Security Classification) ON THE EQUIVALENCE OF LOGICAL DATABASES			
12. PERSONAL AUTHOR(S) Gabriel M. Kuper, Jeffrey D. Ullman, and Moshe Y. Vardi*			
13a. TYPE OF REPORT Technical	13b. TIME COVERED FROM TO	14. DATE OF REPORT (Yr., Mo., Day) 1984	15. PAGE COUNT 8
16. SUPPLEMENTARY NOTATION *IBM Research Laboratory, San Jose CA 95193.			
17. COSATI CODES		18. SUBJECT TERMS (Continue on reverse if necessary and identify by block number) 1984	
FIELD	GROUP	SUB GR.	
19. ABSTRACT (Continue on reverse if necessary and identify by block number) The investigators suggest a new approach to database updates, in which a database is treated as a collection of theories. They investigate two issues: 1) equivalence of databases under update operations, b) simultaneous multiple update operations.			
20. DISTRIBUTION/AVAILABILITY OF ABSTRACT UNCLASSIFIED/UNLIMITED <input checked="" type="checkbox"/> SAME AS RPT. <input type="checkbox"/> DTIC USERS <input type="checkbox"/>		21. ABSTRACT SECURITY CLASSIFICATION UNCLASSIFIED	
22a. NAME OF RESPONSIBLE INDIVIDUAL Dr. Robert N. Buchal		22b. TELEPHONE NUMBER (Include Area Code) (202) 767- 4939	22c. OFFICE SYMBOL NM

AD-A141 541

DTIC FILE COPY

DTIC
ELECTED
MAY 24 1984
A

On the Equivalence of Logical Databases

GABRIEL M. KUPER†

Stanford University
Stanford, California.

JEFFREY D. ULLMAN†

Stanford University
Stanford, California.

MOSHE Y. VARDI‡

IBM Research Laboratory
San Jose, California

Abstract

We suggest a new approach to database updates, in which a database is treated as a collection of theories. We investigate two issues: a) equivalence of databases under update operations, b) simultaneous multiple update operations.

1 Introduction

One of the main problems in database theory is the problem of view updating, i.e., how to translate an update on a user view into an update of the database ([BS], [CA], [DB], [J], [Kc], [KI], [O]). The problem is that in general there is no

unique database update corresponding to the view update. Another problem with updating databases is the problem of updating a database which must satisfy certain integrity constraints ([NY], [T]). The problem here is that the database after the update may no longer satisfy the constraints, in which case we may have to modify other things in the database, to ensure that the integrity constraints still hold. As in the case of view updates, there is not necessarily a unique way to modify the database so that the constraints still hold.

It is shown in [FUV] that even in the absence of constraints, the semantics of updates on the database itself is still not completely clear. They suggest, as in ([Ko], [R]), that the appropriate framework for studying the semantics of updates is to treat the database as a consistent set of sentences in first-order logic, i.e., a theory. Every model of the theory is a possible state of the world. The database is then assumed to be a model of these sentences. Unlike first-order theories, however, the facts that are in the set have a greater significance than those that are just implied by them, and when updating the database we try to keep them as long as possible. When

† Work supported by AFOSR grant 80-0212

‡ This work was done while this author was at Stanford University and supported by a Weizmann Fellowship and AFOSR grant 80-0212.



Accession For	
GRA&I	<input checked="" type="checkbox"/>
TAB	<input type="checkbox"/>
Unannounced	<input type="checkbox"/>
Justification	
By	
Distribution/	
Availability Codes	
Dist	Avail and/or Special
A-1	

Approved for public release;
distribution unlimited.

84 05 15 318

deleting a fact, we try to delete as little as possible, in order to get a theory that does not imply the deleted fact. Similarly, when inserting a new fact, we try to delete as few facts as possible, so that we can then insert the new fact without creating an inconsistent theory.

One problem with this is that there may be several possible different results to an update with no reasonable way to choose between them. One approach to this is suggested in [FUV]. In this approach, the result of the update is defined to be the disjunction of all the possible theories. Two difficulties with this approach are that it requires us to have sentences of a rather complicated syntax, e.g., disjunctions of tuples in a relational database, and that the number of sentences in the database may grow exponentially with each update. In this paper we suggest another approach that consists of defining the result of an update to be a *flock*, which is a collection of theories. We then assume that all we know at each stage about the database is that it is a model of at least one of these theories. With this approach, the sentences we get are of no greater complexity than those that were in the database or those that were inserted, and the number of sentences does not grow as fast as before. It turns out that the two approaches yield different results for the same updates.

In this paper, after presenting the flock approach to updates, we investigate the question when are two flocks equivalent. Since in our approach there is a difference between a fact being in the database and being implied by it, two flocks that are logically equivalent may not be equivalent under all updates. That is, they may have the same models, but there may be updates that when applied to both flocks yield nonequivalent flocks. In this paper, we give necessary and sufficient conditions for equivalence forever and give

several results about batch operations, i.e., inserting or deleting several sentences at the same time.

2 Updates of Theories.

We are going to talk about theories consisting of first-order sentences, that we assume are neither inconsistent nor valid. We shall use the letters S and T to denote theories (i.e., consistent sets of sentences).

Semantics of updates are defined in [FUV] as follows.

Definition 1. (a) A theory T accomplishes the deletion of σ from S if $T \not\models \sigma$. (b) A theory T accomplishes the insertion of σ into S if $\sigma \in T$.

Definition 2. If T_1 , T_2 , and T are theories, we say that T_1 has fewer insertions than T_2 with respect to T iff $T_1 - T \subset T_2 - T$; T_1 has fewer deletions than T_2 with respect to T iff $T - T_1 \subset T - T_2$; and T_1 has fewer changes than T_2 with respect to T iff T_1 has fewer deletions than T_2 , or T_1 and T_2 have the same deletions ($T - T_1 = T - T_2$) and T_1 has fewer insertions than T_2 .

Definition 3. A theory T accomplishes an update u of S minimally iff T accomplishes u and there is no theory T' that accomplishes u and has fewer changes than T with respect to S .

It is shown in [FUV] that T accomplishes the deletion of σ from S minimally iff T is a maximal subset of S that is consistent with $\neg\sigma$, and that $T \cup \{\sigma\}$ accomplishes the insertion of σ into S minimally iff T is a maximal subset of S that is consistent with σ . Nevertheless, there could be many theories that accomplish an update minimally. Suppose that T_1, \dots, T_n are the theories

that accomplish an update u of S minimally. It is argued in [FUV] that the result of u should be a theory T such that

$$\text{Mod}(T) = \bigcup_{1 \leq i \leq n} \text{Mod}(T_i),$$

where $\text{Mod}(S)$ is the set of models of the theory S .

Definition 4. Let T_1, \dots, T_n be theories. The disjunction of these theories is defined to be the theory

$$\bigvee_{1 \leq i \leq n} T_i = \{\tau_1 \vee \dots \vee \tau_n \mid \tau_i \in T_i, 1 \leq i \leq n\}.$$

It is shown in [FUV] that

$$\text{Mod}\left(\bigvee_{1 \leq i \leq n} T_i\right) = \bigcup_{1 \leq i \leq n} \text{Mod}(T_i).$$

Thus they suggest that if T_1, \dots, T_n are the theories that accomplish an update u minimally, then the result of u should be $\bigvee_{1 \leq i \leq n} T_i$.

3 Flocks.

In this paper, we shall talk about another approach to updates, namely using collections of theories. We call these collections *flocks*. The intuitive idea is that since we have many possible theories that accomplish an update minimally, we reflect this ambiguity by keeping all these theories.

Definition 5. A flock S is a set of theories. The models of S are

$$\text{Mod}(S) = \bigcup_{S \in S} \text{Mod}(S).$$

To update a flock we have to update each theory in the flock. Formally:

Definition 6. Let $S = \{S_1, \dots, S_n\}$ be a flock. A flock $T = \{T_1, \dots, T_n\}$ accomplishes an update

u of S minimally if T_i accomplishes the update of S_i minimally, for $1 \leq i \leq n$.

Again, there could be many flocks that accomplish an update minimally. Suppose that T_1, \dots, T_n are the flocks that accomplish an update u of S minimally. As in [FUV], we contend that the result of u should be a flock T such that

$$\text{Mod}(T) = \bigcup_{1 \leq i \leq n} \text{Mod}(T_i).$$

It is easy to show that the flock $\bigcup_{1 \leq i \leq n} T_i$ has this property. This motivates the following definition:

Definition 7. Let S be a flock, and let S_1, \dots, S_n be the flocks that accomplish an update u of S minimally. Then the result of u is the flock $\bigcup_{1 \leq i \leq n} S_i$.

Lemma 1. Let $S = \{S_1, \dots, S_n\}$ be a flock. For each theory S_i , let $S_i^1, \dots, S_i^{j_i}$ be the theories that accomplish the update u of S_i minimally. Then the result of applying u to S is the flock

$$S' = \{S_i^k \mid 1 \leq i \leq n, 1 \leq k \leq j_i\}. \quad \blacksquare$$

In other words, to update a flock, consider each theory in the flock in turn. Take all theories that accomplish the update minimally and put them into the new flock.

Note that if a flock is a singleton, i.e., contains exactly one theory, its models as a theory and as a flock are the same. Also, the flock we get after applying an update to such a flock has the same models as the theory we get by applying the same update under the FUV-approach. However their behavior under future updates may differ, as the following example shows.

Example 1. If we start with the flock $\{\{A, B\}\}$, and delete $A \wedge B$ from it using the flocks approach, we take all the maximal subtheories of

$\{A, B\}$ that do not imply $A \wedge B$, namely $\{A\}$ and $\{B\}$. That is, the resulting flock is $\{\{A\}, \{B\}\}$. If we were then to delete A , and finally delete B , we would end up with the flock containing only the empty theory, i.e., anything is a model of the result. On the other hand, if we were to start with the theory $\{A, B\}$, and delete $A \wedge B$ using the approach of [FUV], we would get the theory $\{A \vee B\}$. This has the same models as the flock $\{\{A\}, \{B\}\}$. However if we now delete A and then delete B , we would end up with $\{A \vee B\}$, which does not have the same models as the empty theory. ■

4 Equivalence Forever.

We have defined the models of a flock to be the union of the models of all the theories in the flock. Therefore, two singleton flocks are logically equivalent if they have the same models. However, this does not guarantee that they will continue to have the same models after any sequence of insertions and deletions, as the next example shows.

Example 2. The two flocks

$$\{\{A, B\}\}$$

and

$$\{\{A, B, A \vee B\}\}$$

are logically equivalent. However, if we delete first A and then B from them (using the flocks rule) we get the nonequivalent flocks $\{\emptyset\}$ and $\{\emptyset, \{A \vee B\}\}$.

We say that two flocks are *equivalent forever* if after applying any sequence of updates we always get two flocks that have the same models. We would like to know when two flocks are equivalent forever. We do not have, at present, a simple necessary and sufficient condition for general flocks. However, for singleton flocks, i.e., flocks

that contain only one theory, equivalence forever can be characterized by the property of covering.

Definition 8. We say that a theory S covers a theory T , iff every sentence τ in T is logically equivalent to a conjunction $\sigma_1 \wedge \dots \wedge \sigma_n$ of sentences in S .

Theorem 1. Let $S = \{S\}$ and $T = \{T\}$ be singleton flocks. Then S and T are equivalent forever if and only if S covers T and T covers S .

Proof: (a) Sufficiency. Assume that S covers T . We show by induction on the number of insertion/deletion steps that we always have

$$(\forall S' \in S')(\exists T' \in T')(S' \text{ covers } T' \wedge T' \text{ covers } S'), \quad (1)$$

where S' and T' are the flocks we get from S and T by doing some insertions and deletions. By our assumption, (1) holds at the beginning, when both flocks are singletons.

Assume that (1) holds after some insertions and deletions. We have to show that (1) continues to hold after deleting a sentence σ . The argument for insertion is similar. We shall use S^1 and T^1 for the flocks before the deletion, S^2 and T^2 for the flocks afterwards.

Let S^2 be any theory in the flock S^2 . We first show that there is some theory in T^2 that covers S^2 .

By the definition of deletion, S^2 must be a maximal subset of some theory S^1 in the flock S^1 that does not imply σ . By the inductive hypothesis, there is a theory T^1 in the flock T^1 such that S^1 covers T^1 and T^1 covers S^1 . Let σ_i be any sentence in the theory S^2 . Since S^2 is a subset of S^1 and T^1 covers S^1 , there are sentences $\tau_{i1}, \dots, \tau_{im_i}$ in T^1 such that $\sigma_i = \tau_{i1} \wedge \dots \wedge \tau_{im_i}$.

Let A be the set of all these τ_{ij} 's, for all σ_i 's in S^2 . We claim that A does not imply σ , the sen-

tence we deleted. Assume otherwise, i.e., $A \models \sigma$. Since each σ_i in S^2 implies all the corresponding τ_{ij} 's in A , we have $S^2 \models A$, and therefore have $S^2 \models \sigma$, a contradiction. Therefore A does not imply σ and can be extended to a maximal subset of T^1 with this property. Call the maximal subset T^2 . Since A covers S^2 , T^2 also covers S^2 . We shall now show that S^2 covers T^2 , thus completing the proof.

Let τ be any sentence in T^2 . We have to show that it is logically equivalent to a conjunction of sentences in S^2 . Since S^1 covers T^1 and T^2 is a subset of T^1 , there are $\sigma_1, \dots, \sigma_k$ in S^1 such that

$$\tau \equiv \sigma_1 \wedge \dots \wedge \sigma_k. \quad (2)$$

Since T^2 covers S^2 , $T^2 \models S^2$. We also know, from (2), that $T^2 \models \tau \models \sigma_i$, for each σ_i . If some σ_i was not in S^2 , the fact that S^2 is a maximal subset of S^1 not implying σ would entail that $S^2 \cup \{\sigma_i\}$ implies σ . But then $T^2 \models \sigma$, a contradiction. This shows that each σ_i is in S^2 , and therefore S^2 covers T^2 .

Now let M be a model of some theory S' in the flock S' . By (1), there is some theory T' in the flock T' , such that S' covers T' . This implies that M is also a model of T' , and therefore M is a model of some theory in the flock T' . By analogous argument we can show that

$$(\forall T' \in T')(\exists S' \in S')(S' \text{ covers } T' \wedge T' \text{ covers } S'),$$

where S' and T' are the flocks we get from S and T by doing some insertions and deletions. Consequently, every model of T' is also a model of S' .

(b) Necessity. Assume that S does not cover T . Then there is a sentence τ in T that is not logically equivalent to a conjunction of any collection of sentences of S .

Let D be the set of sentences in $S \cup T$ that are

not implied by τ . Let R be the set of maximal disjunctions of sentences in D , i.e., the set R of all disjunctions of sentences in D such that if we add any other sentence in D to the disjunction, the result is implied by τ . Formally, R consists of sentences of the form

$$\sigma_1 \vee \dots \vee \sigma_k,$$

where each σ_i is in D ,

$$\tau \not\models \sigma_1 \vee \dots \vee \sigma_k,$$

and if σ is any sentence in D distinct from all the σ_i 's, then

$$\tau \models \sigma_1 \vee \dots \vee \sigma_k \vee \sigma.$$

We show that if we delete the sentences in R from the flock $S = \{S\}$, one by one, in any order, the resulting flock S' will be equal to $\{S - D\}$, and similarly deleting R from $T = \{T\}$ will result in $\{T - D\}$. We will prove this for S . Since the proof will not make use of the fact that τ is not covered by S , the proof will also work for T .

Since no sentence in D is implied by τ , every σ in D can be extended to a maximal disjunction $\sigma \vee \sigma_2 \vee \dots \vee \sigma_k$ that is in the set R . After deleting this disjunction, we get a flock of theories, none of which can contain any of the sentences $\sigma, \sigma_2, \dots, \sigma_k$. Therefore, after deleting all of the sentences in R from S , we get a flock S' of theories, each of which must be a subset of $S - D$.

We now show by induction on the number of deletions, that the result is a singleton flock, consisting of one theory that is a superset of $S - D$. The basis for the induction is the initial flock S . We show by induction, that if we have a flock consisting of one theory that is a superset of $S - D$ and a subset of S , and we delete a sentence in the set R from it, we get a singleton flock, also consisting of one theory that is a superset of $S - D$ and a subset of S .

Suppose that we have such a flock consisting of the theory S' and we delete from it a sentence $\sigma_1 \vee \dots \vee \sigma_k$ from R . Let $D = \{\sigma_1, \dots, \sigma_k\} = \{\alpha_1, \dots, \alpha_m\}$. Since $\sigma_1 \vee \dots \vee \sigma_k$ is maximal, τ implies $\sigma_1 \vee \dots \vee \sigma_k \vee \alpha_i$ for $1 \leq i \leq m$. If we assume that $\{\tau, \alpha_1, \dots, \alpha_m\}$ implies $\sigma_1 \vee \dots \vee \sigma_k$, then it follows that τ implies $\sigma_1 \vee \dots \vee \sigma_k$ contradiction. Consequently, $S' - D \cup \{\alpha_1, \dots, \alpha_m\}$ does not imply $\sigma_1 \vee \dots \vee \sigma_k$. That is the result of deleting $\sigma_1 \vee \dots \vee \sigma_k$ from $\{S'\}$ is $\{S' - \{\sigma_1, \dots, \sigma_k\}\}$. This completes the induction and shows that T' is a singleton, consisting of one theory that is a superset of $S - D$.

By the definition of D , we have $\tau \models S - D$, and therefore τ implies the conjunction of all the sentences in $S - D$. Since τ is not a conjunction of any collection of sentences in S , $S - D \not\models \tau$. Therefore, there must be a model M of $S - D$ that is not a model of τ . Then M is a model of S' . However, since τ is in $T - D$, M is not a model of T' , and the flocks are not forever equivalent. ■

From the proof of this theorem, we immediately get a sufficient condition for equivalence forever of arbitrary flocks.

Corollary. Let S and T be two flocks that satisfy the conditions

$$(\forall S \in \Sigma)(\exists T \in \mathcal{T})(S \text{ covers } T \wedge T \text{ covers } S)$$

and

$$(\forall T \in \mathcal{T})(\exists S \in \Sigma)(T \text{ covers } S \wedge S \text{ covers } T)$$

Then S and T are equivalent forever. ■

Example 3. The flocks $\{\{A, B, A \wedge B\}\}$ and $\{\{A, B\}\}$ are equivalent forever. The flocks $\{\{A, B, A \vee B\}\}$ and $\{\{A, B\}\}$ are not equivalent forever. If we delete A and then B , we get $\{\{A \vee B\}\}$ from the first flock and $\{\emptyset\}$ from the second one.

The converse to the Corollary does not hold, as the following example shows.

Example 4. The two flocks

$$\{\{A, B, A \equiv B\}, \{A, A \equiv B\}, \{B, A \equiv B\}\}$$

and

$$\{\{A, A \equiv B\}, \{B, A \equiv B\}\}$$

are equivalent forever, but do not satisfy the covering condition of the Corollary.

5 Batch Operations.

Batch operations consist of deleting or inserting several sentences at the same time.

Definition 9. Let S be a theory, and Σ a set of sentences. We say that S' accomplishes the deletion of Σ from S if, for each $\sigma \in \Sigma$, $S' \not\models \sigma$. We say that S' accomplishes the insertion of Σ into S iff $\Sigma \subseteq S'$. We say that S' accomplishes an update u of S minimally if S' accomplishes u and there is no theory that accomplishes u with fewer changes.

The above definition is non-constructive in the sense that it does not give us any clue as to how to find those theories that accomplish an update minimally. The following theorem gives a constructive equivalent condition, which generalizes a result of [FUV].

Theorem 2. Let S and T be theories, and Σ a set of sentences. Then

- (1) S accomplishes the deletion of Σ from T minimally if and only if S is a maximal subset of T that is consistent with $\neg\sigma$ for all σ in Σ .
- (2) $S \cup \Sigma$ accomplishes the insertion of Σ into T minimally if and only if S is a maximal subset of T that is consistent with Σ . ■

Definitions 6 and 7 now define batch operations for flocks. Namely, to update a flock, consider each theory in the flock in turn. Take all theories that accomplish the update this theory minimally and put them into the new flock.

The following example shows that show that the batch deletion of Σ does not give the same result as deleting the sentences in Σ one by one.

Example 5. Let S be the singleton flock

$$\{\{A, B, A \equiv B\}\},$$

and let Σ be the set $\{A, B\}$. The result of deleting Σ from S is the flock $\{\{A \equiv B\}\}$. However the result of deleting first A and then B (or first B and then A) is the flock $\{\{A \equiv B\}, \emptyset\}$.

In the case of singleton flocks, the following theorem holds.

Theorem 3. Let $S = \{S\}$ be a singleton flock. If τ , the result of deleting the set of sentences Σ from S is a singleton flock $\{T\}$, then there is a sequence of deletions of single sentences which when applied to S results in T .

Proof: This is proved by replacing the deletion of Σ by a sequence of deletions of maximal disjunctions of elements of Σ , in a similar way to the proof of Theorem 1. ■

Remark. This theorem does not hold if S or T is not a singleton. If we delete the set $\{A, B\}$ from the flock

$$\{\{A \vee B, A \vee \neg B, \neg A \vee B\}\}$$

we get the flock

$$\{\{A \vee B\}, \{A \vee \neg B, \neg A \vee B\}\}$$

which cannot be obtained from a singleton flock by any sequence of insertions and deletions of single sentences. ■

For batch insertions we have the following theorem.

Theorem 4. Let S be a flock, and let $\Sigma = \{\sigma_1, \dots, \sigma_n\}$ a consistent set of sentences. Then the result of inserting Σ into S is the same as first deleting $\neg(\sigma_1 \wedge \dots \wedge \sigma_n)$, and then inserting the σ_i 's, one by one into the result. ■

Proof: This follows immediately from the fact that, for any theory S , S is consistent with Σ iff $S \not\models \neg(\sigma_1 \wedge \dots \wedge \sigma_n)$, and from the fact that, once we have deleted $\neg(\sigma_1 \wedge \dots \wedge \sigma_n)$, insertion of each σ_i consists of simply adding the σ_i to each theory in the flock, without deleting anything else. ■

References

- [BS] Bancilhon, F. and N. Spyrtos, "Update Semantics of Relational Views," ACM Trans. on Database Systems 6 (1981), pp. 557-575.
- [CA] Carlson, C. R. and A. K. Arora, "The Updatability of Relational Views based on Functional Dependencies," Proc. IEEE COMPSAC, 1979, pp. 415-420.
- [DB] Dayal, U. and P. A. Bernstein, "Translation of Update Operations on Relational Views," ACM Trans. on Database Systems 8 (1982), pp. 381-416.
- [FUV] Fagin, R., J. D. Ullman and M. Y. Vardi, "On the Semantics of Updates in Databases," Proc. 2nd ACM Symp. on the Principles of Database Systems, 1983, pp. 352-365.
- [J] Jacobs, B. E., "Application of Database Logic to the View Update Problem," Technical Report TR 960, Dept. of Computer

Science, University of Maryland at College Park, 1980.

- [Ke] Keller, A. M., "Updates to Relational Databases through Views Involving Joins," In *Improving database usability and responsiveness* (P. Scheuermann, ed.), Academic Press, pp. 363-384.
- [Kl] Klug, A. C., "Theory of Database Mapping," Technical Report CSRG 98, Dept. of Computer Science, U. of Toronto, 1978.
- [Ko] Kowalski, R., "Logic as Database Language," Unpublished Manuscript, Dept. of Computing, Imperial College, London, 1981.
- [NY] Nicolas, J. M. and K. Yazdanian, "Integrity Checking in Deductive Databases," In *Logic and Databases* (H. Gallaire and J. Minker, eds.), Plenum Press, 1978, pp. 325-344.
- [O] Osman, I. M., "Updating Defined Relations," *Nat'l Computer Conf.*, Vol. 48, AFIP Press, 1979, pp. 733-740.
- [R] Reiter, R., "Towards a Logical Reconstruction of Relational Database Theory," Unpublished Manuscript, University of British Columbia, 1981.
- [T] Todd, S., "Automatic Constraint Maintenance and Updating Defined Relations," *Proc. IFIP 77* (B. Gilchrist, ed.), North-Holland, 1977, pp. 145-148.